

BACKGROUND OF THE INVENTION

5 The invention relates to the control and/or access of objects, in particular objects known as beans. Beans, for example JavaBeans components (JavaBeans is a trademark of Sun Microsystems, Inc.), are reusable software components which can be manipulated visually in a builder tool (e.g. an editor or graphical user interface builder (GUI builder)). An example of a builder tool is the JavaBeans Development Kit. Further details about beans can be found in many different works, for example in a book entitled Mastering JavaBeans by Lawrence Vanhelsuwé published by Sybex 10 (ISBN 0-7821-2097-0). This is just one example of many broadly equivalent books on the market having "JavaBeans" in the title and describing the JavaBeans. Many of these works, including the book Mastering JavaBeans, supply the Bean Development Kit mentioned above.

15 Beans vary in functionality, but they typically share certain common defining features providing a set of properties, a set of methods for performing actions, and support for events and for introspection, also known as reflection. The properties allow beans to be manipulated programmatically and support customization of the bean. The methods implement the properties. The support for events enables beans to fire events and define the events which can be fired. The support for introspection 20 enables the properties, events and methods of the bean to be inspected from externally.

25 However, beans can generally only be controlled and manipulated within the virtual machine environment in which the beans exist. US patent 5,315,703 and US patent 5,367,633 describe object-based systems where change notification functions are provided. However, these patents describe stand-alone system.

30 It would be desirable to be able to access and control beans remotely, for example over a network. However, this is not been possible without actually including appropriate methods within the beans themselves to expose the methods required. However, it is not desirable to have to specifically adapt beans to permit remote access. It would be desirable to be able remotely to access, control and modify any beans remotely without pre-modification.

NOT CHECKED FOR

002290-1625960

✓ for

SUN REF: P2651

Accordingly it is an object of the invention to address this problem.

002290"46250960

SUMMARY OF THE INVENTION

5 In accordance with an aspect of the invention, there is provided a method of accessing from a client station a target object at a remote station via a telecommunications network, the method comprising steps of:

- a) generating a client object forming a representation of the target object, which client object is configured to identify methods of the target object which are accessible remotely and to implement the remotely accessible methods;
- b) registering the target object and a network adaptor for a network protocol with
10 a framework at the remote station;
- c) associating the client object with a network adaptor for the network protocol at the client machine; and
- d) enabling a client application to access the target object by instantiating the client object.

15 The use of a network protocol adaptor at the client and remote stations and the client object representation of the target object enables remote access, control and modification, if required, of the target object at the remote station without needing to pre-modify the target object or to require the target object to include its own internal remote access methods.

20 In a preferred embodiment of the invention, the target object is associated with the network adaptor by registering the target object with a repository service of a network agent framework, the network adaptor being responsive to remote access requests to cause the framework to initiate access to the target object by referencing the repository service.

25 Preferably, step (a) comprises compiling a target object and generating a client object comprising a target object interface identifying which methods of the target object are accessible remotely and a target object stub implementing the remotely accessible methods. The compiling step can be effected at the remote or the client station and the resulting code be transferred, for example via the network. However,
30 as an alternative, the compiling can be done at both of the client and remote stations.

Step (d) preferably comprises selectively replacing the target object stub for

002290"46250960

dynamically modifying the behaviour of the client application at runtime. For example, it is possible to specify either read-only access or to specify read-write access to the target object using first and second stubs, respectively.

5 The invention finds particular application to objects in the form of beans which comprise a set of properties, a set of methods for performing actions, and support for events and for introspection. The support for introspection enables a compiler to generate the client object(s) and to enable a user application to interact with the client object(s) to access and display at least selected methods of the target object. Indeed, step (a) preferably comprises extracting target object methods by introspection.

10 The invention finds particular, but not exclusive application to a network management system where the target object is a managed object and the client application is a network management application.

15 In accordance with another aspect of the invention, there is provided a remote access support mechanism at a client station for accessing a target object at a remote station via a telecommunications network, the remote access support mechanism comprising:

a client object forming a representation of the target object, which client object identifies methods of the target object which are accessible remotely and implements the remotely accessible methods; and

20 a network adaptor responsive to the client object,
wherein the client object is configured to be instantiated by a client application for enabling the client application to access and modify the target object.

25 In accordance with a further aspect of the invention, there is provided a remote access support mechanism at a first machine permitting remote access from a client machine to a target object at the first machine via a telecommunications network, the remote access support mechanism comprising:

at least one target object;

at least one network adaptor supporting a network protocol:

30 the target object(s) and the network adaptor(s) being registerable with a framework at the first machine and the network adaptor being responsive to remote access requests from the client machine in accordance with the protocol to access the

002290" 46250960

SUN REF: P2651

target object via the framework.

In accordance with yet another aspect of the invention, there is provided a remote access support mechanism on a data carrier for a computer for permitting remote access from a client machine to a target object at the first machine via a telecommunications network, the remote access support mechanism comprising:

at least one target object;

at least one network adaptor supporting a network protocol:

the target object(s) and the network adaptor(s) being registerable with a framework at the first machine and the network adaptor being responsive to remote access requests from the client machine in accordance with the protocol to access the target object via the framework.

002290"46250960

BRIEF DESCRIPTION OF THE DRAWINGS

Exemplary embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference signs relate to like elements and in which:

Figure 1 is a schematic representation of three stations connected via a telecommunications network;

Figures 2 and 2A form a schematic representation of a computer server for a station of Figure 1;

Figure 3 is a schematic representation of an agent for a managed station;

Figure 4 is an alternative representation of the agent of Figure 3;

Figure 5 is an example of a configuration for an agent;

Figure 6 is a flow diagram illustrating operations using an agent as shown in Figure 5;

Figure 7 is an example of a configuration of a management system;

Figure 8 is another example of a configuration of a management system;

Figure 9 illustrates an aspect of the creation of the configuration of Figure 8;

Figure 10 is a flow diagram illustrating operations of the system of Figure 8;

Figure 11 is a flow diagram illustrating the operation of a naming service;

Figure 12 is a flow diagram illustrating the operation of the creation of a generic agent;

Figures 13 and 14 illustrate alternative operations of a compiler;

Figures 15A and 15B are used to illustrate the effect of compiling a management bean;

Figures 16A and 16B are also used to illustrate the effect of compiling a management bean; and

Figure 17 is a flow diagram illustrating steps in generating a management system.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 is a schematic representation of a multi-station network based system 1 with three stations, or nodes, or machines 3, 4 and 5 connected via a network 2.

5 The network 2 can be based on a public switched telephone network and/or a local area network and/or a dedicated network connecting computer and other equipment within a local area and/or over a wider area and/or an open network such as the Internet or an intranet, or combination thereof or indeed any other type of telecommunications network which can support the exchange of telecommunications management information. The network can have any desired structure, such as a level structure, a pyramidal hierarchy, etc.

10 Any one or more of the stations 3, 4 or 5 can be configured as a network management station. In the example shown in Figure 1, it is assumed that station 3 is a management station and stations 4 and 5 are managed stations. It will be appreciated that any number of management stations and managed stations may be provided and that the three stations of Figure 1 are for illustrative purposes only. Also, the managing and managed functions could be interchanged or indeed both managing and managed functions could be supported at one and the same station.

15 The stations can take on many forms. For the purposes of illustration it is assumed that both comprise computer workstations. Figure 2 is a schematic representation of the configuration of a computer workstation. As illustrated in Figure 2, this is implemented by a server computer 10 comprising a system unit 11 with a display 8, keyboard and other input devices 9. Figure 2A is a schematic block representation of aspects of the contents of the system unit 11. As illustrated in Figure 2A, the system unit includes a processor 17, memory 18, magnetic and/or optical disk drives 13 and 14, and a communications adaptor 16 for connection to one or more telecommunication lines 15 for connection to the telecommunications network 2. As illustrated in Figure 2A, the components of the system unit are connected via a bus arrangement 19. It will be appreciated that Figures 2/2A are a general schematic representation of one possible configuration for a server computer for forming a router. It will be appreciated that many alternative configurations could

00000000000000000000000000000000

be provided.

Where the workstation implements a management station, management application or applications and interface structures are typically provided by software which is stored in the memory of the workstation and executed by the processor of
5 the workstation.

Where the workstation implements a managed station, an agent is responsive to remote management requests via the network from the management applications to provide an interface to managed objects at the managed stations. The agent and managed object structures are typically provided by software which is stored in the
10 memory of the workstation and executed by the processor of the workstation.

The objects typically comprise parameters and methods used to model a piece of equipment, a component of that equipment, an operation of the equipment or a component or resource thereof, and so on.

The management of a telecommunications network requires applications of
15 various sizes and complexities. Heavy managers, middle managers, extensible agents, smart agents and appliances all have a role in the management of a telecommunications network. A network management system incorporating the present invention provides an extensible agent framework allowing all of these different application types to be built on the same architecture. This extensible agent
20 framework is provided by a component of the network management system. Alternative terms for the extensible agent framework in this context could be a "dynamic framework" or "open framework" or "runtime component", although the terms "framework" or "extensible agent framework" will be used herein.

The network management system is supplied with a set of core management
25 services. A choice can be made from this set of services, and it can be extended to develop specific applications. Different services are loaded statically or dynamically into the framework to meet the requirements of a particular application.

A managed object in an example of a network agent incorporating the present invention is preferably implemented as a bean, more preferably a JavaBeans
30 component. A bean (and a JavaBeans component) is a reusable software component which can be manipulated visually in a builder tool (e.g. an editor or graphical user

002290-46250960

interface builder (GUI builder). An example of a builder tool is the JavaBeans Development Kit. Beans vary in functionality, but they typically share certain common defining features providing a set of properties, a set of methods for performing actions, and support for events and for introspection. The properties
5 allow beans to be manipulated programmatically and support customization of the bean. The methods implement the properties. The support for events enables beans to fire events and define the events which can be fired. The support for introspection enables the properties, events and methods of the bean to be inspected from externally. Operations such as GET, SET, ACTION, CREATE and DELETE can
10 be supported.

A managed object in the agent is manageable as soon as it is registered with the framework. This arrangement enables an agent developed in accordance with this network management system to be manageable with minimal impact on the design of the agent.

15 As indicated above, an example of the network management system uses the JavaBeans component model, thereby easing the development of applications. In this example, all of the core management services are provided as JavaBeans components. Thus access can be had to them using a Java application builder, such as the well known JavaBeans Development Kit. Managed objects are developed as JavaBeans
20 components. A JavaBeans component in the network management system agent can be accessed locally or remotely. This means that when developing managed objects with the network management system, it is not necessary to know to which communications protocol the managed object will be accessed.

25 The network management system simplifies the development of extensible agents. A set of core management services that the network management system provides can be extended and loaded into an agent while it is running. Most of the core management services are optional. This means that an agent developed using the network management system need only implement the service it uses. This feature enables the development of agents of differing sizes and complexities.

30 Agents developed using the network management system are also smart agents. A smart agent provides the services needed to process management requests. In a

0022901625060

SUN REF: P2651

smart agent, much of the processing can be done locally in the agent itself, reducing the load on the network connection between the agent and managing system.

Figure 3 illustrates an aspect of the architecture of a network management system agent, including the relationships between the components of the network management system. Figure 3 also shows the relationship between the network management system agent and management applications.

As shown in Figure 3, the network management system agent consists of a number of components inside a Java virtual machine (VM). These include:

- m-beans 29;
- the framework 24;
- core management services 25, 26, 27, 28;
- managed objects adaptor servers 30, 32, 34, 36 and 38.

These components will be described in more detail below.

A **managed object** is a software abstraction of a resource that is controlled and monitored by an agent. A managed object (eg. 28) is referred to as a management bean or m-bean. In an example of the network management system, all m-beans are implemented as JavaBeans components. Therefore, these can be accessed using a conventional Java application builder, such as the JavaBeans Development Kit mentioned earlier.

As for any other managed object, an m-bean has a set of properties, can perform a set of actions, and can emit a set of notifications or events. The network management system enables a distinction to be made between a read-only and a read-write property in an m-bean.

An m-bean is manageable as soon as it is registered with the framework 24. When an m-bean is registered, an object name is associated with it. The object name uniquely identifies the m-bean within the m-bean repository (see below). It enables a management application to identify the m-bean on which it is to perform a management operation. The object name of an m-bean is an arbitrary name that does not depend in any way on how the m-bean is implemented.

The framework 24 controls the management services and m-beans of an agent. The framework, or runtime component, 24 is

SUN REF: P2651

5 a JavaBeans component which comprises a set of properties, a set of methods for performing actions, and support for events and for introspection. The properties include getter and setter properties. The methods include methods to implement the getter and setter properties. The framework is also able to effect add object and remove object functions.

10 Whenever an agent 20 is requested to perform a management operation received through a network adaptor, the framework 24 calls the appropriate service to perform the requested operation. The framework 24 also handles communications between m-beans 28 and the managed object adaptor servers 30-38. An m-bean can query the framework 24 to obtain information on other m-beans 28 loaded into the same instance of the framework 24. Only one instance of the framework 24 is permitted within a virtual machine 22.

15 The network management system provides a number of core management services. In an example of the system, the core management services are defined as Java interfaces. The core management services are optional. This means that an agent developed using the network management system need only implement the services it uses. The core management services can be registered as m-beans, which allows some management operations to be formed on them for tuning their behaviour.

20 In the preferred example of the system, the core management services are provided as JavaBeans components. It is therefore possible to access them using a conventional Java application builder, such as the JavaBeans Development Kit mentioned earlier.

A number of the core management services are now to be described.

25 The **m-bean repository service 27** obtains pointers to m-beans. Each time an m-bean is registered with the framework 24, the framework 24 calls the m-bean repository service 27 to store the identity of the m-bean. A name is associated with an m-bean. When querying the m-bean repository 27, agent and manager applications can identify an m-bean by its name. Different implementations of the m-bean repository service 27 are possible. One implementation uses a relational database for storing persistent m-beans, whereas another implementation employs simple memory for storing non-persistent m-beans.

30

002290-46250960

Bearing in mind the structure provided by the repository unit, an alternative representation of the relationship between components of the agent is illustrated in Figure 4. This takes account of the registration of the beans for the managed object adaptor servers, the management services and the managed objects with the m-bean repository service bean.

A **filtering service** 29 selects m-beans to be the subject of a management operation. Selection is based on the presence and values of specific attributes. For example, a filter could select all the m-beans for which the attribute color is set to red.

A **metadata service** 25 provides information on the structure of an m-bean. For example, the framework 24 queries the metadata service 25 to access methods for getting and setting up a property within an m-bean. The metadata service 25 is based on the Reflection API provided by the JavaBeans Development Kit for performing introspection.

A **dynamic class loading service** loads new classes into the framework 24. A new class is loaded when the remote entry requests the creation of instances of a class that is not loaded into the framework 24. A new class can be loaded from a remote class server. The dynamic loading service also enables core management services to be added to a network management system agent while it is running. For example, an agent could be started without a filtering service. Then, later on, the filtering service could be added dynamically to the agent when it is required.

An **access control service** can be provided to control access to m-beans. Before attempting to perform a management operation on an m-bean, the framework 24 can be arranged to query the access control service to verify that the operation is valid.

An **event service** can be provided to receive event reports from m-beans and to forward them to any entity that has requested to receive them.

A **relationship service** can be provided to enable relationships between m-beans to be defined when they are required. The relationships do not need to be defined in advance. Information on the relationships between m-beans is not stored with the m-beans themselves, but can be obtained from the relationship service.

5 A **dynamic native library loading service** can be provided for loading native code into the framework 24. A native library can be loaded when a new class that includes native code is loaded. The library loaded depends on the hardware platform and operating system on which the framework is running. The native library can be loaded from a remote entity.

There now follows a description of the managed object adaptor servers 30-38.

10 A **managed object adaptor server** is a protocol adaptor that provides an abstraction of a communications protocol. Each managed object adaptor server provides access to m-beans through a particular communications protocol. A managed object adaptor server enables management applications to perform management operations on a network management system agent.

15 For a network management system agent to be manageable, it is connected via at least one managed object adaptor server. However, a network management system agent can be connected to any number of managed object adaptor servers, allowing it to be queried by remote management applications that use different protocols.

The network management system provides managed object adaptor servers for standard and proprietary protocols.

For example, managed object adaptor servers can be provided for one or more of the standard protocols: HTML/HTTP; SNMP; and CORBA, by way of example.

20 The managed object adaptor servers for standard protocols communicate directly with the management applications that use them.

25 For example, an HTML/HTTP managed object adaptor server enables a user to use a web browser to view management information contained in m-beans and to perform management operations on a network management system agent. The HTTP/HTML managed object adaptor server obtains management information from m-beans and generates HTML pages representing this management information.

An SNMP managed object adaptor server can be arranged to use a specially defined SNMP management information base (MIB) to enable the SNMP manager to perform management operations on a network management system agent.

30 The network management system can also be arranged to provide managed object adaptor servers for one or more of the following proprietary protocols: RMI;

002290-46250960

SUN REF: P2651

Java/HTTP; and Secure Sockets Layer (SSL). In one example of the network management system, a managed object adaptor client offers a Java API. Accordingly, any management application that uses a managed object adaptor client will be written in Java language.

5 Agents developed using the network management system can be managed using different communications or management protocols. To be managed using a standard management protocol, a network management system agent needs to be connected to the managed object adaptor server for that protocol. The managed object adaptor service for standard protocols communicates directly with the
10 management applications that use them.

An example of this structure, using the agent representation of Figure 4, is shown in Figure 5, where the agent is accessed by means of a web browser 46. Here an HTML adaptor allows beans to be mapped to an HTML page. The use of a protocol such as HTML enables the browser 46 at the client station 90 to browse
15 beans at the remote station 20 and access and where necessary modify them using conventional web browser functions. In accordance with the structure shown in Figure 4, the repository service 27 is registered with the framework 24, and the HTML adaptor 34 and the bean(s) 29 are registered with the repository service, whereby the bean(s) are effectively registered with an HTML page supported by the
20 HTML adaptor 34. In Figure 5 like reference numerals relate to like features of the arrangement shown in Figure 4.

Figure 6 is a flow diagram illustrating steps for enabling the modification of properties of the beans in the remote machine. In step 92, the HTML adaptor 34 is initiated in the virtual machine 22 at the remote station 20 and in step 94 the bean(s)
25 29 of that virtual machine which are to be accessible remotely are registered with the framework, or more particularly with the repository service 27 as described above. This means that when the HTML adaptor queries the framework 24, the framework 24, with reference to the repository service, is able to identify the beans 29 to be accessed and to permit access thereto by the HTML adaptor.

30 The HTML adaptor 34 allows communication over the network using conventional HTTP exchanges. It behaves as an HTTP server. When it receives a

SUN REF: P2651

request, it dynamically generates a page containing a list of beans (objects) 29 currently registered with the repository object 27.

A bean is represented in HTML as an HTML table wherein:

- a first column contains a property name;
- 5 a second column contains a property type;
- a third column contains the access right (read/write);
- a fourth column contains a property value.

As mentioned above, if the property is read/write, an HTML form is generated.

10 In step 96, the beans are displayed at the client station (represented by display 98) using the HTML representation of a bean as described above by accessing the HTML page using a conventional web browser which communicates with the HTML adaptor using HTTP exchanges. The user is able then to select, at step 100, a bean using conventional web browser functions. The web browser will then issue an
15 HTTP GET request to the HTML adaptor 34. The HTML adaptor employs introspection to extract the bean properties and then returns an HTML post response to the browser, whereby the browser may display the properties, and possibly also the actions and events supported by the bean, as represented at 102. By further use of the browser using conventional browser functions, the user is able to select and
20 specify modifications to aspects of the bean, for example changes to the properties. By a further exchange of HTML GET and/or SET requests and POST responses, the web browser and the HTML adaptor are able to modify, at step 104, the corresponding properties of the bean at the remote station and to display these changes to the user at the client station.

25 Thus, this mechanism provides a computer-implemented method of accessing from a client machine an object such as a bean at a remote machine via a telecommunications network by mapping the object to an externally accessible machine page at the remote machine and browsing the object via the machine page using a browser.

30 Another example of the structure shown in Figure 3, this time using the representation of Figure 3, is shown in Figure 7, where a network management

SUN REF: P2651

system agent 20 is managed by an SNMP manager application. In Figure 7 like reference numerals relate to like features of the arrangement shown in Figure 3.

5 An example of a network management system represented in Figure 8 provides an adaptor API to enable Java management applications to communicate with the network management system agents. The adaptor API provides managed object adaptor clients for accessing managed objects through a particular communications protocol. The network management system defines a representation of m-beans for Java management applications and provides a compiling tool for generating such a representation automatically from an m-bean. A name service is supplied to allow
10 Java management applications to be independent of a particular communications protocol.

A managed object adaptor client is an abstract Java class that enables Java management applications to access managed objects. The programmatic representation of the managed object to the Java management application is
15 determined by the managed object adaptor client. Such a mechanism allows different representations of the same managed object to be presented to different Java management applications. The network management system provides managed object adaptor clients for accessing managed objects through one or more of the following protocols: RMI; HTTP; and SSL.

20 The managed object adaptor clients provide a definition of an adaptor managed object interface. The adaptor managed object interface enables Java management applications to perform one or more of the following management operations on a network management system agent:

- retrieve m-beans;
- 25 - get or set the properties of remote m-beans;
- call methods of remote m-beans;
- create instances of m-beans;
- delete m-beans; and
- receive events emitted by remote m-beans.

30 A managed object adaptor client provides a Java management application with "handles" on managed objects in a remote agent. These handles enable the Java

002290-46250960

management application to manipulate the managed objects directly. The Java management application does not need information on the protocol used by the managed object. All the Java management application needs is the class of object that the managed object represents. For example, a Java management application for
5 handling accounts uses an abstract class for representing accounts. To manipulate an account, the Java management application obtains an account managed object from the managed object adaptor client. It then casts the account managed object into the abstract class that represents the account. In this way, the application code is independent of how the managed object is implemented.

10 Figure 8 is a schematic representation of the relationship between a client bean (c-bean) 54 and an m-bean 28. A c-bean 54 is a representation of an m-bean to a Java management application. In the preferred embodiment of the invention, a c-bean 54, like an m-bean 28, is implemented as a JavaBeans component. A c-bean 54 defines how a Java management application accesses an m-bean 28.

15 As seen in Figure 8, a c-bean 54 comprises a managed object interface (MO interface) 56 which defines which of the methods of an m-bean are accessible to a Java management application, and a managed object stub (MOSub) 58 that implements the methods defined in the MO interface 56.

A Java management application obtains a reference to a c-bean by using an
20 adaptor MO interface 60. The adaptor MO interface instantiates the c-bean 54. The same implementation of a c-bean 54 can run on any managed object adaptor client that implements the adaptor MO interface 60. Different implementations of the same managed object can be presented to different Java management applications. Therefore, a single m-bean can be associated with several c-beans 54.

25 A Java management application performs management operations on an m-bean by calling methods of its associated c-bean 54. To the Java management application, a c-bean 54 is a local representation of the remote Java object (an m-bean 28).

The adaptor MO interface 60 instantiates the c-bean 54. The same
30 implementation of a c-bean can run on any managed object adaptor client 62 which implements the adaptor MO interface 60. Different representations of the same

managed object can be presented to different Java management applications. Thus, a single m-bean can be associated with several c-beans.

A Java management application performs management operations on a m-bean by calling methods of its associated c-bean. To the Java management application, a c-bean is a local representation of a remote Java object (an m-bean).

Figure 9 is a schematic representation of the generation of a c-bean 54 from an m-bean. In an embodiment of the invention a c-bean is generated automatically from an m-bean 28 using the Reflection API. The generated c-beans exhibit the same properties, methods and events as the m-beans. This is the case, for example, where no access control policy is enforced.

In order to provide the automatic generation of a c-bean 54 from an m-bean 28, a compiler 60 takes as an input the m-bean 28 and generates as an output the MO interface and MO stubs of a c-bean 54. For example, when an m-bean 28 representing a Java class named *account* is compiled, the compiler 60 generates an MO interface 56 named *accountMO* and a Java class named *accountMOSTUB* 58, which implements the *accountMO* interface 56.

A Java management application can be developed using the MO interface definitions. By loading different stubs, the *adaptorMO* interface can modify the behaviour of the Java management application at run time. For example, if read-only stubs are loaded, the Java management application will not be able to modify the properties in the m-bean 28.

The compiler 60 can generate read-only or read-write stubs. The generated stubs make use of the *adaptorMO* interface. Therefore, their behaviour is the same on any managed object adaptor client that implements the *adaptorMO* interface, regardless of the implementation of the managed object adaptor client.

Figure 10 is a block diagram of steps for accessing a bean at a remote (server) station from a local management (client) station using a structure as shown in Figure 8.

The management application (e.g. a Java management application) at the management station generates, at step 80, a get request to the adaptorMO at the client station. The adaptorMO, with the Mostub and the network adaptor at the client

SUN REF: P2651

station generate, at 81, a request in accordance with an appropriate network protocol (e.g. HTTP).

5 Thus the request sent via the network to the managed system could, for example, be in the form of a GET request for a management property of a managed object.

The appropriate managed object adaptor server 30-38, depending on the protocol used, receives the external request at step 82. It will then access, at step 83, the framework 24 to get an appropriate method. The framework gets, at step 83, a managed object method for the request and returns, at step 84, the management
10 property of the managed object to the adaptor, which in turn returns, at step 85, composes a return message with the result in accordance with the same network protocol (e.g. HTTP). The result message is received, at step 86, at the client adaptor and adaptorMO, which returns the result to the management application.

A name service is provided which allows management applications to be
15 independent of a particular communications protocol. Java management applications use the name service to determine which managed object adaptor client to use for accessing an agent. Figure 11 is a schematic flow diagram illustrating the operation of the name service.

In step 62, the management application passes the identity of the agent to be
20 accessed to the main service.

In step 64, the name service returns the class name of the managed object adaptor client, for example *sunw.jaw.moa.rmi* in the case of an agent access through the Java RMI system.

In step 66, the Java management application uses its default class loader to
25 dynamically load the managed object adaptor client and instantiate it. The Java management application can then interact with the agent through the managed object adaptor client, regardless of the communications protocol.

As mentioned above, a management bean, or m-bean is a managed object in
30 a network management system agent. A managed object is a software abstraction of a resource that is controlled and monitored by the agent. In an example of a network management system, all m-beans are implemented as JavaBeans components. They

002290-46250960

SUN REF: P2651

can be accessed using a conventional Java application builder, such as the Java Beans Development Kit. Within an m-bean, it is possible to call and use services provided by the network management system.

5 A JavaBeans component includes properties which form discrete, named attributes which can affect the appearance or the behaviour of the JavaBeans component. For example, an m-bean representing an Ethernet driver might have a property named *Ipackets* that represents the number of incoming packets. Properties can have arbitrary values, including both built-in Java end types and class or interface types such as *Java.awt.color*. Properties are always accessed via method calls on the
10 object that owns them. For readable properties there is a *get* method to read the property value. For writeable properties, there is a *set* method to allow the property value to be updated.

A default design pattern can be used for locating properties:

```
public PropertyType get PropertyName();  
15 public void set PropertyName (PropertyType value);
```

If a class definition contains a matching pair of *get PropertyName* and *set PropertyName* methods with the return type of the getter corresponding to the parameter type of the setter, then these methods define a read-write property. If a class definition contains only one of these methods, the name defines either a read-
20 only or write-only property called *PropertyName*.

In addition for Boolean properties, it is possible to define a *get* method using the following design pattern:

```
public boolean isPropertyName();
```

The *isPropertyName* method may be provided instead of a *getPropertyName*
25 method where it may be provided in addition to a *getPropertyName* method.

An index property is an array *PropertyElement*[], that is accessed by methods of the form:

```
public PropertyElement getPropertyName (int index);  
public void setPropertyName (int index, PropertyElement b).
```

30 If a class definition contains either kind of method, *PropertyName* is an index property. These methods can be used to read and write a property value. These

SUN REF: P2651

methods can be defined in addition to the methods defined for simple properties. Therefore, an index property can be represented by four accessor methods.

By default, the following design pattern is used to determine which events an m-bean can multicast:

```
5      public void addEventListenerType(EventListenerType a);  
      public removeEventListenerType(EventListenerType a);
```

Both of these methods take the *EventListenerType* type argument, where the *EventListenerType* type extends the *java.util.EventListener* interface, where the first method starts with *add*, the second method starts with *remove*, and where the
10 *EventListenerType* type name ends with *Listener*.

This design patent assumes that the Java bean is acting as a multicast event source for the event specified in the *EventListenerType* interface.

To conform to the JavaBeans model, all public methods of a JavaBeans component should be exposed as external methods within the component environment
15 for access by the components. By default, this includes property accessor methods, and in the event listener registry method.

In addition to the JavaBeans component model for design pattern elements, the network management system can define an action as a public method of an m-bean that makes sense to be called remotely. Action facilitates the differentiation of an m-
20 bean public method which is exposed for other local m-beans from public methods that can be called remotely. The design pattern for an action is as follows:

```
      public AnyJavaType performAnAction (AnySignature);
```

M-beans can contain native libraries, that is a library not written in the language (e.g. Java) of the m-bean. The network management system can provide
25 a mechanism for loading a native library from the same remote class server as the m-beans. To enable an m-bean to use this mechanism, a static *loadLibrary* method of the Framework class can be called in which the caller includes a reference to the Java class of the caller. Such information is used by the framework 24 for identifying the class loader by which the class is loaded.

30 The core management services mentioned above are for many of the management operations common to all agents, thereby simplifying the development.

of agents. By providing these core services, the network management system enables efforts to be concentrated on developing those parts of an agent which are specific to a particular application, namely the m-beans and the applications to control and manage them.

5 Figure 12 is a schematic flow diagram of the initialisation of a network management system agent 20. The initialisation process comprises:

- in step 70, creating an instance of the framework 24;
- in step 72, adding the m-bean repository service 27;
- in step 74, adding the metadata service 29; and
- 10 - in step 76, adding at least one managed object adaptor server (30-38) so that the agent can be access by management applications.

Once the network management system agent 20 has been initialised, no further management services need to be added before an agent is started. These can be added dynamically to the agent while it is running.

15 The framework 24, controls the management services and m-beans of an agent 20 developed using the network management system. In the preferred embodiment it is implemented by the Java class `java.jaw.agent.cmf.Framework`. An agent must contain one instance of the framework, that is, in the preferred embodiment, one instance of the `java.jaw.agent.cmf.Framework` class.

20 In the preferred embodiment, m-beans can be managed only if they are registered with an object name in the m-bean repository 27 of the agent 20. Accordingly, the m-bean repository service 27 is added in step 72 before the agent 20 becomes operational. The m-bean repository service 27 is used for storing and retrieving the association between an m-bean and its object name. The m-bean repository service of the preferred embodiment is defined as the Java interface `java.jaw.agent.services.MoRepSrvIf`. An agent can only be associated with the one m-bean repository service at any one time. However, it is possible to change the m-bean repository service with which the agent is associated while the agent is running.

25 The m-bean repository can be implemented as a volatile storage or as a persistent storage. In a volatile repository, all the information on m-beans is stored in memory. All of the information in a volatile repository is lost when the agent is

SUN REF: P2651

stopped. Accordingly, when an agent is started, with a volatile repository, it has to reload the information in the repository. With a persistent repository, all the information on m-beans is stored in a database whereby the information in a persistent repository is not lost when the agent is stopped. It is also possible to implement a
5 mixed repository whereby information on m-beans is stored in memory or in a database.

The metadata service referenced above is used to obtain the properties and actions supported by an m-bean. A preferred implementation of the metadata service is based on the Reflection API provided by the known Java Development Kit for
10 performing introspection.

As mentioned above, at least one managed object adaptor service should be running in the server virtual machine as the network management system agent. The network management system does not require a managed object adaptor server to conform to a specific interface definition or implementation. The managed object
15 adaptor server is arranged to access the framework 24 to retrieve and change information contained in the agent. The managed object adaptor servers provided are implemented as m-beans. Examples of managed object adaptor servers have been described above. In the preferred embodiment of the invention, the managed object adaptor servers are implemented as appropriate Java classes.

For example, an RMI managed object adaptor server can be implemented as
20 a Java class `sunw.jaw.agent.adaptor.rmi.AdaptorServerImpl`. It enables Java management application to access an agent using the Java remote method invocation (RMI) system. As described above, a Java management application accesses this server through a managed object adaptor client implemented as the Java class
25 `sunw.jaw.agent.adaptor.rmi.AdaptorClient`.

Core management services can be added to an agent while it is running, once the agent has been initialised as described above. It is possible to add a core service to an agent in either of the following ways:

- directly calling a set method for the service within the Framework class;
- 30 - adding the service to the m-bean repository in the same way as for an m-bean.

SUN REF: P2651

Adding a core management service directly gives faster performance than adding a core management service to the m-bean repository. This is because the framework does not need to query the m-bean repository in order to obtain the service. However, certain restrictions can apply to a core management service that has been added directly:

- the service is not visible to remote applications; and
- it is not possible to store information on the service in persistent storage.

Accordingly, where it is desirable for a core management service to be visible to remote applications, it is necessary to add the service to the m-bean repository. If it is desired to store information on a core management service in persistent storage, it is also necessary to add the service to the m-bean repository. The m-bean repository to which the service is added, must support persistent storage.

A Class Service Name contains names by which the framework identifies the services that are implemented for an agent. The framework retrieves the service it requires as follows:

1. The framework checks if a service has been defined using a direct set method. If a service has been defined in this way, the framework uses this service.
2. If the service has not been defined using a direct set method, the framework queries the m-bean repository to obtain all the m-beans that are instances of the class that implements the service. For example, for the metadata service, the framework queries the m-bean repository to obtain all the instances of the class named `ServiceName.META`.

- if the repository contains several instances of the class, the framework uses the first instance returned by the m-bean repository.
- if the m-bean repository contains no instances of the class, the framework throws a `ServiceNotFound` Exception.

Various operations can be performed in a network management service agent. For example, an object in an agent can use core management services for:

- instantiating m-bean;
- registering m-beans with the m-bean repository;

SUN REF: P2651

- retrieving m-beans from the m-bean repository;
- getting and setting the values of properties within m-beans; and
- defining relationships between m-beans.

5 An object name uniquely identifies an m-bean. Management applications use object names to identify the m-beans on which to perform management operations. Any naming scheme could be used. For example, in the preferred embodiment of the invention, a naming scheme defined by Microsoft Corporation for the Hyper Media Management Scheme (HMMS) could be used.

10 In order to instantiate an m-bean, one of the following methods of the Framework class can be called:

- newObject to use default storage mechanism for storing the m-bean;
- newDBObject to specify the m-bean is persistent.

With either of these methods, it is necessary to provide:

- the Java class of the m-bean to be instantiated; and
- 15 - the object name to be used for registering the m-bean.

By default, the framework 24 uses the default class loader to locate the Java class of the m-bean to be created. It then creates an instance of the class. Once the m-bean has been instantiated, it is initialised and registered so that it is accessible to the framework 24. It is possible to initialise and register an m-bean by using:

- 20 - a method defined in the m-bean itself; or
- the framework 24.

For initialising a register of an m-bean using a method defined in the m-bean itself, the Java class definition of the m-bean should contain:

- an initialisation method;
- 25 - the code required to enable the m-bean to register itself with the m-bean repository.

Once the m-bean has been instantiated, the framework 24 uses the metadata service 27 to find the initialisation method in the newly created m-bean. If such a method is present in the m-bean, the framework 24 calls it giving:

- 30 - a reference to itself as a first parameter;
- the object name for use in registering the m-bean as a second parameter.

SUN REF: P2651

The m-bean is therefore able to register itself with the m-bean repository using the code provided.

If an m-bean is not provided with the initialisation method, the framework initialises and registers the m-bean using functions provided for this purpose.

5 Registering a JavaBeans component with the m-bean repository 25 enables the component to be managed by the agent 20. Registering the JavaBeans component does not require modification of code within the JavaBeans component itself. Instead, all that is required is the addition of code for registering it in the m-bean repository. Therefore, it is possible to register any existing JavaBeans component in the m-bean repository. Once registered, the agent 20 manages the JavaBeans component in the same way as any m-bean. When an m-bean is registered, it is assigned an object name. An object name can be explicitly specified. If an object name is not explicitly specified, the framework 24 assigns a default name to the m-bean.

15 The network management system provides services for retrieving m-beans from the m-bean repository. These services enable the retrieval of m-beans using:

- pattern matching on the object name; or
- queries (filters) on the Java properties they contain.

By using pattern matching on the object names of m-beans, it is possible to retrieve:

- 20
- a specific m-bean using its full object name;
 - a set of m-beans sharing the same logical class as expressed in the object name;
 - a set of m-beans sharing the same domain name; or
 - all the m-beans contained in an agent.

25 Using queries enables the retrieval of m-beans according to Java properties and their values within m-beans. The m-bean repository evaluates properties if it is able to do so. Otherwise the framework evaluates queries itself. To determine whether a repository is able to assess queries, the framework causes a query method for this purpose.

30 The network management system provides services for getting and setting properties of m-beans. If the agent provides a metadata service, in a call to the get

SUN REF: P2651

or set method of the m-bean, all that needs to be supplied are:

- the name of the property to be retrieved or set;
- the object name of the m-bean that contains the property.

If the agent does not provide a metadata service, it is still possible to call
5 directly to the get or set method of an m-bean. In this case it is also necessary to
supply to the caller the name and signature of the method to call.

A relationship service enables relationships between m-beans to be defined
when they are required. The relationships do not need to be defined in advance.
Information on the relationships between m-beans is not stored with the m-beans
10 themselves, but is stored with the relationship. A relationship needs to be registered
with the m-bean repository. A relationship is defined by:

- a set of roles, for example in an ownership relationship a person owns a book
and a book is owned by a person;
- a degree which corresponds to the number of required roles in a relationship.

15 The m-beans involved in a relationship are referred to in a relationship by
their object names. A specific class loader can be added to the agent at start-up or
while the agent is running. It is possible to have several class loaders within the
same agent, provided that they are all registered with the m-bean repository. When
the creation of a new object is requested, it is possible to specify the class loader to
20 be used for loading the object. In this case the class loader is identified by its object
name. The system provides several implementations of a network class loader, each
implementation using a different protocol and requiring a different class server.

The system provides event handling services for filtering, logging and
forwarding events through different communications protocols. To emit an event, a
25 send event method is called. When this method is called, the framework 24 retrieves
all the m-beans corresponding to an event handling service and calls an event handling
method of each event handler retrieved. This method is responsible for handling the
events.

Figure 9 is a schematic representation of the compiling of a c-bean 54 from
30 an m-bean 28. The compiler 60 implements one translation scheme for generating
c-beans. It is, however, possible to implement different translation schemes.

depending on the requirements.

Figure 13 illustrates an example of the output of the compiler 60 for compiling a single m-bean called A. Figure 14 shows the output of the compiler 60 if a compiled m-bean includes *aListener* for a specific event *AnEvent*.

5 Thus, if the m-bean contains listeners, the compiler 60 generates:

- a Java interface for the listener to be included in the MO interface;
- a listener stub which is an implementation of the m-bean listener for catching m-bean events and forwarding them to the framework 24; and
- a Java management application's view of the event associated to the listener

10 referred to as *EventMO*.

The compiler 60 parses an m-bean using the applicable design parameters. After parsing, the compiler 60 uses a number of rules for generating the c-bean.

Each property of the m-bean will be present in the c-bean with the same accessor methods. So, if a property is read-only in the m-bean, the property will be

15 read-only in the c-bean.

Figure 15A is an illustration of a simple c-bean which is generated when compiling an m-bean defined by a code example shown in Figure 15B. In addition, the compiler 60 generates a file containing an implementation of the Simple MO interface.

20 In addition to the property accessors, the compiler 60 generates code only for public methods for which it makes sense to provide remote access. Other public methods do not appear in the c-bean generated by the compiler 60.

Figure 16A shows a subset for an action method in a c-bean of the MO interface which the compiler 60 will generate when compiling the action method in

25 an m-bean defined as in Figure 16B.

If an m-bean contains a listener called A, the compiler 60 includes a listener called AifMO in the c-bean.

When using the c-bean, an application will have to implement the Aifmo interface to add or remove the listener in the c-bean. Generally, a listener is an

30 interface containing a certain number of methods. Each method has one input parameter. The input parameter extends an event object concerned.

002290-16250960

In an example, each method defined in listener A refers to an event object which, for the purpose of this example is called *AnEvent*. Accordingly, in the Aifmo interface, the event object is called *AnEventMO*. Thus, for a listener A, the compiler 60 generates files:

- 5 - Aifmo.java;
- AnEventMO.java.

In addition, the compiler 60 generates an implementation of listener A named AStub.java.

Codes generated by the compiler 60 complies with the design parameters designed by the JavaBeans component model. Accordingly, the objects generated by the compiler 60 can be integrated in development environments that comply with this model. In addition, the compiler 60 adds some public methods which do not follow the design patterns defined by the JavaBeans component model. The added methods are designed to limit the network traffic between an m-bean and a c-bean. For example, by calling one function on a c-bean, it is possible to read all the properties of the corresponding m-bean.

The compiler 60 generates Java source code. It is possible to edit the generated code and modify it to define a specific view of an m-bean. Instead of modifying both the interface and the stub, it is better to keep the interface and modify only the stub. The code generated by the compiler 60 enables an application to be built using the interface. At one time the behaviour will change depending on which stubs are loaded by the *adaptorMO*. For instance, the compiler can generate read-only or read-write stubs for the same interface. Accordingly, an m-bean browser can be developed based on the interface. As mentioned above, the browser will thereby have a different behaviour depending on whether the read-only or read-write stubs are loaded.

The adaptorMO interface is a Java interface defined for managing the agent 20. The network management system provides several implementations of the adaptorMO interface based on different communications protocols, as described above. However, the adaptorMO interface is protocol independent. Accordingly, any piece of code written using the interface can run on any of the implementations

002290-1625050

provided by the network management system.

5 Within the adaptorMO interface there are two different levels. A low level where remote objects (m-bean) are manipulated using their names, and a high level where remote objects (m-bean) are manipulated using a local view (c-bean) of the remote objects. The high level is built on top of the low level interface.

10 Using the low level interface is practical for building generic applications (such as an HTML object viewer or MIB browser). Using the high level interface is much easier than using the lower level interface. However, it means that the application knows the semantic of the c-beans it manipulates. In addition, it requires the application (or the adaptorMO interface) to have access to MOs and MOStubs. A first step in initialising an adaptor consists of initialising an implementation of the adaptorMO interface.

15 In order to initialise an adaptor, the client applications invokes a "connect" method defined in the adaptorMO interface. Parameters are provided which relate to the host name of the agent, the port number to use and a logical name which is generally dependent on the underlying communication mechanism. The different piece of information could be obtained from a name server or directory service at the same time the implementation name of the adaptor to use is given. Depending on the underlying communication mechanism used by the adaptorMO interface, the call to "connect" might not involve any message exchange between the client and agent.

20 An adaptor makes use of:

- a name server for getting the Java class name for use for representing a specific m-bean (known through its object name);
- a class loader for loading c-beans.

25 If all the Java classes for the c-beans are present on the machine where the client application is running, there is no need to use a specific class loader. Otherwise, it is possible to use a network class loader for getting the classes over the network.

30 For using a network class loader, a client application needs to instantiate the network class loader. When instantiating the object, the application provides an object name. The object name can contain any domain and any class name.

SUN REF: P2651

However, the object name should contain the following key properties:

- host (the host name where the class server is running);
- port (the port number to use);
- service (the name of the RMI service to invoke).

5 Once an adaptor is initialised, the application is ready to perform management operations on the remote agent. An application can retrieve a subset or all of the m-beans managed by an agent. When retrieving objects, an application can specify queries. The results of the retrieval can be obtained under two different schemes:

- a list of object names (represented by a Vector);
- 10 - a list of c-beans (for each object name retrieved the adaptor will instantiate a c-bean).

 In order to read a property of a remote m-bean, the property name is needed when using the low level adaptorMO interface level. When using the high level interface, the c-bean is retrieved and then a getter associated to the property is
15 invoked.

 Setting a property of a remote m-bean, a property name and the property object type is required when using the low level adaptorMO interface level. When setting a value, it is possible to specify the name of an operator class. On the agent side, the specified operator is instantiated and invoked for setting the property value.
20 When using the low level adaptorMO interface level, it is possible to set several properties through one method call using a list of modifications.

 When using the high level adaptorMO interface level, a c-bean is obtained and the developer then invokes the set associated to the property.

 Through the adaptorMO interface, it is possible to request instantiated of m-beans within a remote agent. When requesting the instantiation, it is possible specify
25 a new class loader through which the agent is supposed to load the new class to instantiate. The class loader can be specified using its object name. If no class loader is specified, the agent uses the default class loader. When instantiating a remote m-bean, it is possible to directly obtain a c-bean for representing the newly
30 created m-bean. If no name is provided and if a name server is specified, the adaptor queries the name server in order to obtain the name of the Java class to instantiate on

SUN REF: P2651

the agent's side. Otherwise, it is the responsibility of the agent to determine the class name of the class to instantiate. When instantiating an m-bean in the agent, it is possible to explicitly request the object to be persistent.

5 Through the adaptorMO interface, it is possible to transfer Java objects from the client to the agent. In order to do this, the adaptorMO interface serialises the object, sends the object over and deserialises the object in the agent.

Through the adaptorMO interface, it is also possible to remove an m-bean object from the remote agent. The m-bean is not removed from the virtual machine, but only from the object repository of the agent.

10 Figure 17 is a flow diagram giving an overview of the steps in creating and operating a management system as described above including steps of defining a network management model including at least one management bean using a bean-based environment and compiling said model to implement said computer network management system in said bean-based environment.

15 In step 110, a model is created using a bean-based environment. A preferred bean-based environment is Java environment with the beans being JavaBeans.

As mentioned above, beans provide a set of properties, a set of methods for performing actions, and support for events and for introspection. Conventionally, the properties allow beans to be manipulated programmatically and support customization
20 of the bean, the methods implement the properties and the support for events enables beans to fire events and define the events which can be fired. The support for introspection enables the properties, events and methods of the bean to be inspected from externally.

Accordingly, step 110 includes generating at least one said management bean
25 providing at least one property for modelling a property of a managed resource, and/or a method for modelling an action for said managed resource and/or support for at least one event for modelling an event of said resource and/or support for introspection permitting external analysis of the composition of said bean. This step also includes defining the relationship and interactions between management beans as
30 representative of the relationships and interactions between the managed resources. This step can also include defining at least one listener event in respect of at least one

management bean.

It is recognised for the first time in respect of the present network management system that beans can be used beyond their conventional role to provide a management vehicle (management bean) for directly modelling resources to be managed. For example, a property in a management bean can be used for modelling a property (e.g. a resource attribute such as the size of a memory, the number of messages received in a buffer, etc) of a managed resource. A method of a management bean can be used for modelling an action (e.g. a system rest) of a managed resource. A bean can also be used to provide support for an event (for example a disk full event) for a managed resource. For example, a management bean can provide support for a listener event, whereby one management bean can be responsive to an event on another management bean. By means of the support for introspection a management bean can permit external analysis of the composition of the bean and the exchange of information between beans. Management beans can include definitions of relationships and interactions between management beans as representative of the relationships and interaction between the managed resources.

As beans are reusable software components which can be manipulated visually in a builder tool (e.g. an editor or graphical user interface builder (GUI builder)), in step 110 a user can use a conventional builder tool, such as the JavaBeans Development Kit, for generating the management system model including beans defining system resources, their functions and the relationships between and interactions of the system resources. The bean model is defined within a virtual machine forming the bean-based environment, for example a model using JavaBeans within a Java virtual machine. This greatly facilitates the generation of the management system model. Verification and debugging of the model can be readily performed using introspection and other techniques.

In step 112, once the model has been generated, the model can be compiled directly using a conventional compiler for the bean-based environment, for example the compiler 60 shown in Figure 9. By defining a model in a bean-based environment, it is possible directly to compile the bean-based model to form a bean-based implementation using a standard compiler for the bean-based environment, for

002290-1625060

example by compiling the JavaBeans using a Java compiler. Accordingly, the resulting bean management system is also defined within the same Java environment. This greatly simplifies this stage of the generation of management system as the compiler forces a reliable and homogeneous implementation of the model.

5 At runtime, therefore, in step 114, the management system described earlier in this document provides a robust and readily modifiable structure, without the difficulties and disadvantages of prior network management system.

Step 114 includes the steps described in respect of Figure 12 of registering one or more management bean(s) with the extensible agent framework; registering one or
10 more network adaptor(s) (e.g. network adaptor bean(s)) for a network communications protocol with the extensible agent framework; and enabling external access via the network to the management bean(s) via the network adaptor(s).

As described with respect to Figure 12, the extensible agent framework can include an associated repository bean and the steps of registering one or more
15 management bean(s) and/or network adaptor(s) can comprise registering one or more management bean(s) and/or network adaptor(s) with the repository bean.

Thus there has been described, in particular with reference to Figures 8-10, a method and mechanism for accessing from a client machine or station a target object at a remote machine or station via a telecommunications network.

20 A client object forming a representation of the target object is generated, in the particularly described example by compiling a target object and generating a client object comprising a target object interface identifying which methods of the target object which are accessible remotely and a target object stub implementing the remotely accessible methods. The target object stub can preferably be replaced for
25 dynamically modifying the behaviour of the client application at runtime.

The client object is configured to identify methods of the target object which are accessible remotely and to implement the remotely accessible methods. The target object is associated with a network adaptor for a network protocol at the remote station. The client object is associated with a network adaptor for the network
30 protocol network protocol at the client machine. A client application is thereby able to access the target object by instantiating the client object.

002290-4529960

In the preferred implementation described the client object and the target object are beans, each of which comprises a set of properties, a set of methods for performing actions, and support for events and for introspection, whereby a client bean extracts target object methods by introspection. A preferred application of the invention is in a network management system, whereby the target object is a managed object and the client application is a network management application.

Thus there has been described a method and mechanism whereby remote object (remote bean) access can be provided without needing to modify the target object or to include specific communications code in the target object by generating a client object or objects forming a representation of a target object. The client object(s) identify(ies) methods of the target object which are accessible remotely and implement the remotely accessible methods. The target object is associated with a network adaptor for a network protocol at the remote station. The client object is associated with a network adaptor for the network protocol at the client machine. A client application can then access the target object by instantiating the client object.

It will be appreciated that although particular embodiments of the invention have been described, many modifications/additions and/or substitutions may be made within the spirit and scope of the present invention as defined in the appended claims. With reference to those claims, it is to be noted that combinations of features of the dependent claims other than those explicitly enumerated in the claims may be made with features of other dependent claims and/or independent claims, as appropriate, within the spirit and scope of the present invention.

002290-46250960